CLAIMS

What is claimed is:

1. A method enabling a flow-based data collection scheme, comprising

receiving a flow, the flow comprising at least one packet;

monitoring the flow in relation to at least one flow attribute;

associating a traffic type to the flow;

upon termination of the flow, composing a flow data record comprising a traffic type identifier corresponding to the traffic type associated with the flow and the at least one monitored flow attribute; and

storing the flow data record in a database.

2. The method of claim 1 further comprising

parsing at least one packet associated with the flow into a flow specification, wherein said flow specification contains at least one instance of any one of the following: a protocol family designation, a direction of packet flow designation, a protocol type designation, a pair of hosts, a pair of ports, a pointer to a MIME type, a pointer to an application-specific attribute;

matching the flow specification of the parsing step to a plurality of traffic types, each of the traffic types defined by one or more matching attributes; and thereupon,

having found a matching traffic type in the matching step, associating said flow specification with a traffic type from the plurality of traffic types.

3. The method of claim 1 further comprising

parsing at least one packet associated with the flow into a flow specification, wherein said flow specification contains at least one instance of any one of the following: a protocol family designation, a direction of packet flow designation, a protocol type designation, a pair of hosts, a pair of ports, a pointer to a MIME type, a pointer to an application-specific attribute;

matching the flow specification of the parsing step to a plurality of hierarchically-recognized traffic types represented by a plurality of nodes, each node having a traffic specification defining at least one matching attribute; thereupon,

having found a matching node in the matching step, associating the first flow specification with a traffic type of said plurality of hierarchically-recognized traffic types.

4. The method of claim 1 further comprising

parsing at least one packet associated with the flow into a first flow specification, wherein said first flow specification contains at least one instance of any one of the following: a protocol family designation, a direction of packet flow designation, a protocol type designation, a pair of hosts, a pair of ports, a pointer to a MIME type, a pointer to an application-specific attribute;

matching the first flow specification of the parsing step to a plurality of recognized traffic types represented by a plurality entries in at least one traffic type identification table, each entry in the traffic type identification table including at least one matching attribute; thereupon,

having found a matching traffic type in the matching step, associating the first flow specification with a traffic type of said plurality of recognized traffic types in the at least one traffic identification table.

5. The method of claim 1 wherein the at least one flow attribute is one of any of the following: a first packet time, a last packet time, the number of packets in the flow, the number of bytes in the flow, the number of retransmitted bytes in the flow, or a round trip time.

6. The method of claim 1 wherein the flow data record further includes one of any of the following: source address, destination address, source port number, destination port number, VLAN identifier, type of service identifier, a protocol family designation, a direction of packet flow designation, a protocol type designation, a

protocol message designation, a first packet time, a last packet time, the number of packets in the flow, the number of bytes in the flow, the number of retransmitted bytes in the flow, or a round trip time.

7. The method of claim 1 wherein termination of a flow is determined based on a threshold period of time and the time of the last packet in the flow.

8. The method of claim 1 wherein termination of a flow is determined based on a protocol message.

9. The method of claim 8 wherein the protocol message is a TCP FIN packet.

10. The method of claim 1 wherein at least one mapping exists between a traffic type identifier transmitted in flow data records and a traffic type name; and wherein the method further comprises

periodically storing the at least one mapping in the database.

11. The method of claim 1 wherein at least one mapping exists between a traffic type identifier transmitted in flow data records and a traffic type name; and wherein the method further comprises

periodically transmitting the at least one mapping for storage in the database.

12. The method of claim 11 wherein the at least one mapping is transmitted in one or more mapping messages; and wherein the mapping messages further include time stamps.

13. An apparatus enabling a flow-based data collection scheme, comprising

a packet processor operative to

receive a flow, the flow comprising at least one packet;

associate a traffic type to the flow;

monitor the flow in relation to at least one flow attribute; and

a flow data record emitter operative to:

upon termination of a flow, compose a flow data record comprising a traffic type identifier corresponding to the traffic type associated with the flow and the at least one monitored attribute; and

transmit the flow data record to a data collector.

14. The apparatus of claim 13 further comprising

a traffic classification database operative to

store mappings between traffic type names and traffic type identifiers;

classify a flow into a traffic type based on at least one attribute of the flow; and

wherein the flow data record emitter is further operative to:

transmit the mappings between the traffic type names and traffic type identifiers to the data collector.

15. The apparatus of claim 13 wherein the packet processor is further operative to

parse at least one packet associated with the flow into a flow specification, wherein said flow specification contains at least one instance of any one of the following: a protocol family designation, a direction of packet flow designation, a protocol type designation, a pair of hosts, a pair of ports, a pointer to a MIME type a pointer to an application-specific attribute;

match the flow specification to a plurality of traffic types, each of the traffic types defined by one or more matching attributes; and thereupon,

having found a matching traffic type in the matching step, associate said flow specification with a traffic type from the plurality of traffic types.

16. The apparatus of claim 13 wherein the packet processor is further operative to

parse at least one packet associated with the flow into a flow specification, wherein said flow specification contains at least one instance of any one of the

following: a protocol family designation, a direction of packet flow designation, a protocol type designation, a pair of hosts, a pair of ports, a pointer to a MIME type, a pointer to an application-specific attribute;

match the flow specification to a plurality of hierarchically-recognized traffic types represented by a plurality of nodes, each node having a traffic specification defining at least one matching attribute; thereupon,

having found a matching node in the matching step, associate the first flow specification with a traffic type of said plurality of hierarchically-recognized traffic types.

17. The apparatus of claim 13 wherein the packet processor is further operative to

parse at least one packet associated with the flow into a first flow specification, wherein said first flow specification contains at least one instance of any one of the following: a protocol family designation, a direction of packet flow designation, a protocol type designation, a pair of hosts, a pair of ports, a pointer to a MIME type, a pointer to an application-specific attribute;

match the first flow specification to a plurality of recognized traffic types represented by a plurality entries in at least one traffic type identification table, each entry in the traffic type identification table including at least one matching attribute; thereupon,

having found a matching traffic type in the matching step, associate the first flow specification with a traffic type of said plurality of recognized traffic types in the at least one traffic identification table.

18. The apparatus of claim 13 wherein the at least one flow attribute is one of any of the following: a first packet time, a last packet time, the number of packets in the flow, the number of bytes in the flow, the number of retransmitted bytes in the flow, or a round trip time.

19. The apparatus of claim 13 wherein the flow data record further includes one of any of the following: source address, destination address, source port number, destination port number, VLAN identifier, type of service identifier, a protocol family designation, a direction of packet flow designation, a protocol type designation, a protocol message designation, a first packet time, a last packet time, the number of packets in the flow, the number of bytes in the flow, the number of retransmitted bytes in the flow, or a round trip time.

20. The apparatus of claim 13 wherein termination of a flow is determined based on a threshold period of time and the time of the last packet in the flow.

21. The apparatus of claim 13 wherein termination of a flow is determined based on a protocol message.

22. The apparatus of claim 21 wherein the protocol message is a TCP FIN packet.

23. A system enabling a flow-based data collection scheme, comprising

at least one network device disposed in a communication path between first and second networks; the first network device comprising

a packet processor operative to

receive a flow, the flow comprising at least one packet;

associate a traffic type to the flow;

monitor the flow in relation to at least one flow attribute; and

a flow data record emitter operative to:

upon termination of a flow, compose a flow data record comprising a traffic type identifier corresponding to the traffic type associated with the flow and the at least one monitored attribute; and

transmit the flow data record to a data collector; and

a data collector operative to:

receive flow data records from the first network device; and

store the flow data records in a searchable database.

24. The system of claim 23 further comprising

a traffic classification database operative to

store mappings between traffic type names and traffic type identifiers;

classify a flow into a traffic type based on at least one attribute of the flow; and

wherein the flow data record emitter is further operative to:

transmit the mappings between the traffic type names and traffic type identifiers to the data collector.

25. The system of claim 23 wherein the packet processor is further operative to

parse at least one packet associated with the flow into a flow specification, wherein said flow specification contains at least one instance of any one of the following: a protocol family designation, a direction of packet flow designation, a protocol type designation, a pair of hosts, a pair of ports, a pointer to a MIME type a pointer to an application-specific attribute;

match the flow specification to a plurality of traffic types, each of the traffic types defined by one or more matching attributes; and thereupon,

having found a matching traffic type in the matching step, associate said flow specification with a traffic type from the plurality of traffic types.

26. The system of claim 23 wherein the packet processor is further operative to

parse at least one packet associated with the flow into a flow specification, wherein said flow specification contains at least one instance of any one of the following: a protocol family designation, a direction of packet flow designation, a protocol type designation, a pair of hosts, a pair of ports, a pointer to a MIME type, a pointer to an application-specific attribute;

match the flow specification to a plurality of hierarchically-recognized traffic types represented by a plurality of nodes, each node having a traffic specification defining at least one matching attribute; thereupon,

having found a matching node in the matching step, associate the first flow specification with a traffic type of said plurality of hierarchically-recognized traffic types.

27. The system of claim 23 wherein the packet processor is further operative to

parse at least one packet associated with the flow into a first flow specification, wherein said first flow specification contains at least one instance of any one of the following: a protocol family designation, a direction of packet flow designation, a protocol type designation, a pair of hosts, a pair of ports, a pointer to a MIME type, a pointer to an application-specific attribute;

match the first flow specification to a plurality of recognized traffic types represented by a plurality entries in at least one traffic type identification table, each entry in the traffic type identification table including at least one matching attribute; thereupon,

having found a matching traffic type in the matching step, associate the first flow specification with a traffic type of said plurality of recognized traffic types in the at least one traffic identification table.

28. The system of claim 23 wherein the at least one flow attribute is one of any of the following: a first packet time, a last packet time, the number of packets in the flow, the number of bytes in the flow, the number of retransmitted bytes in the flow, or a round trip time.

29. The system of claim 23 wherein the flow data record further includes one of any of the following: source address, destination address, source port number, destination port number, VLAN identifier, type of service identifier, a protocol family designation, a direction of packet flow designation, a protocol type designation, a

protocol message designation, a first packet time, a last packet time, the number of packets in the flow, the number of bytes in the flow, the number of retransmitted bytes in the flow, or a round trip time.

30. The system of claim 23 wherein termination of a flow is determined based on a threshold period of time and the time of the last packet in the flow.

31. The system of claim 23 wherein termination of a flow is determined based on a protocol message.

32. The system of claim 31 wherein the protocol message is a TCP FIN packet.